




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁵ : H04L 12/56	A2	(11) International Publication Number: WO 91/14326 (43) International Publication Date: 19 September 1991 (19.09.91)
(21) International Application Number: PCT/US91/01513 (22) International Filing Date: 5 March 1991 (05.03.91) (30) Priority data: 488,693 5 March 1990 (05.03.90) US (60) Parent Application or Grant (63) Related by Continuation US 488,693 (CIP) Filed on 5 March 1990 (05.03.90) (71) Applicant (for all designated States except US): MASSACHUSETTS INSTITUTE OF TECHNOLOGY [US/US]; 77 Massachusetts Avenue, Cambridge, MA 02139 (US).	(72) Inventors; and (75) Inventors/Applicants (for US only): ARORA, Sanjeev [IN/US]; 69 Chestnut Street, Cambridge, MA 02139 (US). LEIGHTON, Frank, Thomson [US/US]; 965 Dedham Street, Newton Center, MA 02159 (US). MAGGS, Bruce, M. [US/US]; 60 Wadsworth Street, Apt. 21-C, Cambridge, MA 02142 (US). (74) Agents: SMITH, James, M. et al.; Hamilton, Brook, Smith & Reynolds, Two Militia Drive, Lexington, MA 02173 (US). (81) Designated States: AT (European patent), AU, BE (European patent), CA, CH (European patent), DE (European patent), DK (European patent), ES (European patent), FI, FR (European patent), GB (European patent), GR (European patent), HU, IT (European patent), JP, KR, LU (European patent), NL (European patent), NO, SE (European patent), US. Published Without international search report and to be republished upon receipt of that report.	

(54) Title: SWITCHING NETWORKS WITH EXPANSIVE AND/OR DISPERSIVE LOGICAL CLUSTERS FOR MESSAGE ROUTING



(57) Abstract

A class of switching networks is comprised of expansive logical clusters and/or dispersive logical clusters. These clusters are of low degree. The class of networks include multibutterfly networks as well as multi-Benes networks. These networks provide for fault tolerance and routing and for efficient routing. Moreover, routing is provided in a non-blocking fashion.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	ES	Spain	MG	Madagascar
AU	Australia	FI	Finland	ML	Mali
BB	Barbados	FR	France	MN	Mongolia
BE	Belgium	GA	Gabon	MR	Mauritania
BF	Burkina Faso	GB	United Kingdom	MW	Malawi
BG	Bulgaria	GN	Guinea	NL	Netherlands
BJ	Benin	GR	Greece	NO	Norway
BR	Brazil	HU	Hungary	PL	Poland
CA	Canada	IT	Italy	RO	Romania
CF	Central African Republic	JP	Japan	SD	Sudan
CG	Congo	KP	Democratic People's Republic of Korea	SE	Sweden
CH	Switzerland	KR	Republic of Korea	SN	Senegal
CI	Côte d'Ivoire	LI	Liechtenstein	SU	Soviet Union
CM	Cameroon	LK	Sri Lanka	TD	Chad
CS	Czechoslovakia	LU	Luxembourg	TG	Togo
DE	Germany	MC	Monaco	US	United States of America
DK	Denmark				

- 1 -

SWITCHING NETWORKS WITH EXPANSIVE AND/OR DISPERSIVE
LOGICAL CLUSTERS FOR MESSAGE ROUTING

Related Patent Applications

05 This application is a Continuation-in-Part of
pending United States Patent Application Serial No.
07/488,693, filed March 5, 1990, entitled "Switching
Networks with Expansive and/or Dispersive Logical
Clusters for Message Routing".

Background of the Invention

10 A switching network typically is made of input
ports and output ports that are interconnected by
switches and wires. The switching network serves
primarily to correctly route messages from the input
ports to the output ports. Each wire in the network
15 serves as a conduit for transmitting a message from
one of its ends to the other of its ends. The term
wire, in this context, or the terms connection or
connector, includes any means for communicating data
between switches, such as electrical wires, parallel
20 groups of wires, optical fibers, multiplexed channels
over single wires, or free space radio or optical
communication paths.

The switch is an atomic unit that resembles a
switching network in function (i.e., a switch has
25 inputs and outputs and connects the inputs to the
outputs in any desired pattern). The degree of a
switch is the number of inputs and outputs in the
switch. For example, as shown in Figure 1 a 2x2
switch 2 has a degree of four.

- 2 -

A switching network may route any kind of digital or analog data including voice or video signals. It can also route address information that specifies the correct output for the message, and
05 routing information that helps direct the message to the correct output or that establishes communications links such as in a telephone network. In some networks, the routing is accomplished by setting switches so that input ports become directly
10 connected to output ports (e.g., in a telephone network). In other networks, the inputs ports do not become directly connected to the output ports. Instead, the messages are routed as packets through the network in steps.

15 Switching networks are widely used to route messages and to establish communications among multiple parties. Typical examples of networks in which switching networks are used include telephone networks, data networks, computer networks, and
20 interconnection networks in parallel data processing systems.

There are several varieties of switching networks that are classified by the manner in which messages are handled by the network. Common types of
25 switching networks include packet-switching, circuit-switching, cut-through, and worm-hole networks.

In a packet-switching network, packets are treated as atomic objects. At each time step, each
30 wire in the packet-switching network can transmit an entire packet from one switch to another switch. If necessary, packets may be queued in buffers located

- 3 -

at the switches or at the wires of the network. These networks are also often referred to as store-and-forward networks. The name "store and forward" is derived from the characteristic of the networks that packets are temporarily stored in the queues and then forwarded to the next destination in the network.

A circuit switching network is appropriate when messages are too large to be treated as atomic objects (such as packets). In this type of network, a dedicated path is established in the switching network between the sender and receiver of each message. The paths corresponding to different messages are disjoint (i.e., they do not share any wires). Once a path is established, the sender can transmit an arbitrarily long message to the receiver without interference from the other messages. This model is also called the lock-down model, and most closely resembles the approach adopted by current telephone networks.

The wormhole and cut-through networks are best classified as lying in a class situated between the packet and circuit-switching networks. In a wormhole network, a packet is assumed to consist of a sequence of flits (a flit is typically a bit or a byte). At the end of each wire is a buffer that can hold a small number of flits (typically two flits). A packet is not stored entirely in one buffer, but instead is spread out over a quantity of wires indicated by the packet length (number of flits) and by the buffer size. A packet can be thought of as a worm proceeding head-first through the network.

- 4 -

Behind the head, each flit of the worm advances only if there is adequate space in the buffer at the end of the next wire to hold the flit. When the head moves, the buffer space it frees up trickles back to the tail, allowing the entire worm to move. If blocked, a packet compresses (like an accordion) to a length acceptable for the buffer size at the appropriate node. The integrity of the packet is also preserved (i.e., it cannot be cut in half by another packet). In a cut-through network, the buffer size is large enough that the entire packet can accumulate at a single node.

Summary of the Invention

The present invention is comprised of a novel class of switching networks comprised of low-degree, expansive logical clusters and/or low-degree, dispersive logical clusters, and of methods described below for routing messages on these networks in an efficient on-line fashion. The methods for routing messages are superior to previously known methods in that they are fast, fault-tolerant, on-line, and non-blocking. These attractive features are attained by virtue of the expansion and/or dispersion properties of the logical clusters of the switching network.

In accordance with the present invention, a logical cluster comprises a first set and a second set of switches having inputs for receiving messages and outputs for outputting messages. The second set of switches is divided into one or more disjoint groups of switches. The first set of switches and possibly the second set of switches make local

- 5 -

routing decisions. Connectors are provided for connecting the first set of switches to the second set of switches. The connectors interconnect the switches such that an output of each of the first set of switches is connected to an input of a switch in each of the groups of the second set of switches.

The connectors interconnect the first set of switches and the second set of switches so that the logical cluster exhibits an expansion property. In particular, there exists, for every set of k switches in the first set of switches, at least βk switches in each group of the second set of switches are connected to the k outputs of the first set of switches. $\beta > 1$ and $k \leq \alpha N$, where N equals the number of inputs into the first set of switches. Further, $\alpha \leq 1/\beta$. The switches of the logical cluster may each have two inputs and two outputs.

The present invention also envisions a logical cluster that is dispersive. Specifically, for every set of k switches in the first set of switches, there are at least δk switches in each group of the second set of switches that are connected to precisely one of the k switches in the first set of switches, where $k \leq \alpha N$, and N and k are positive integers. Both such logical clusters are ideal for use in multibutterfly switching networks and multi-Benes switching networks. δ and α are positive constants less than one.

The present invention also embodies a multibutterfly switching network made from the merger of individual butterfly switching networks. For purposes of referencing the butterfly switching networks it is helpful to number them 1 through d ,

- 6 -

where d is an integer. Each butterfly switching network has N inputs, and it is made of levels and rows of switches. The butterfly switching networks are merged such that given a set of permutations $(\Pi^1, \dots, \Pi^{(d-1)})$ where $\Pi^k = \langle \pi_0^k, \pi_1^k, \dots, \pi_{lgn}^k \rangle$ and $\pi_L^k: [0, N/2^L - 1] \rightarrow [0, N/2^L - 1]$, a switch in row $jN/2^L + i$ of level L of butterfly switching network k is merged with a switch row $jN/2^L + \pi_L^k(i)$.

Similarly the present invention embodies a multi-Benes switching network that is formed in a manner similar to that of the multibutterfly switching network. In particular, it is formed from d individual Benes switching networks numbered 1 through d . The merger can be described using a set of permutations. Specifically, given a set of permutations $(\Pi^1, \dots, \Pi^{(d-1)})$ where $\Pi^k = \langle \pi_0^k, \pi_1^k, \dots, \pi_{2lgn}^k \rangle$ and where $\pi_L^k: [0, N/2^{lgn-L-1}] \rightarrow [0, N/2^{lgn-L-1}]$ for $0 \leq L \leq lgn$, a switch in row $jN/2^{lgn-L} + i$ of level L of Benes switching network number k is merged with a switch in row $jN/2^{lgn-L} + \pi_L^k(i)$ of level L of Benes switching network number $k+1$ for all $1 \leq k \leq (d-1)$, all $0 \leq i \leq N/2^{lgn-L-1}$, all $0 \leq j \leq 2^{lgn-L-1}$, and all $0 \leq L \leq lgn$. Moreover, for $lgn \leq L \leq 2lgn$, $\pi_L^k: [0, N/2^{L-lgn-1}] \rightarrow [0, N/2^{L-lgn-1}]$, a switch in row $jN/2^{L-lgn} + i$ of level L of Benes switching network number k is merged with a switch in row $jN/2^{L-lgn} + \pi_L^k(i)$ for all $1 \leq k \leq (d-1)$, all $0 \leq i \leq N/2^{L-lgn-1}$, all $0 \leq j \leq 2^{L-lgn-1}$.

Examples of these switching networks (i.e. multibutterfly switching networks and multi-Benes switching networks) are where $d=2$, implying a twin multibutterfly or a 2-multi-Benes. Often times in

-7-

the prior art it is common to select the permutations Π so that they are an identity map. In other words, $i = \pi(i)$. This choice of Π creates a dilated butterfly switching network or a dilated Benes switching
05 network.

Such networks as the multibutterfly switching network and the multi-Benes switching network may be used to produce a very fault tolerant routing scheme. In particular, according to the present invention the
10 output connections of each switch in a switching network organized into levels of switches are examined to determine whether the switch is available. The switch is declared unavailable if the examining step reveals that the switch is faulty or
15 busy, i.e., unusable, or does not have a sufficient quantity of connections to available switches in each output group for each logical cluster of the switching network. Such unavailable switches are avoided in routing messages across the switching
20 networks. To further boost fault tolerance, it is preferred that the input switches of each logical cluster also be inspected to determine how many of the input switches are faulty. Where a number of input switches of a logical cluster that are faulty
25 exceeds a predetermined threshold all of the switches in a logical cluster are declared faulty as well as any descendant switches of the logical cluster. Optimally this additional examining step proceeds from the $\log_2 N$ th level of switches backwards towards
30 the 0 level of switches. In general, the fault tolerance approach may be expanded to heighten the integrity of routing of messages in a switching network. In accordance with this generalized method

- 8 -

all switches are initially declared as available. If a switch is faulty, busy or connected to an insufficient number of properly operating switches it is declared unavailable. Messages are routed
05 exclusively over available switches.

The present invention encompasses packet switching routing strategies as well as circuit switching routing strategies. In accordance with the packet switching method of the present invention,
10 packets of information are routed across a switching network comprised of several levels of switches by first dividing packets of information into waves. Once the packets are divided into waves, the waves are sent from even levels of switches to odd levels
15 of switches during a first time frame. In a second time frame, the waves of packets are sent from odd levels of switches to even levels of switches. In addition, it is preferred that colors are assigned to wires so that each switch is incident to one wire of
20 each color. The colors have a predefined hierarchy. During either of the above sending steps, the packet are sent sequentially over the wires to interconnect the switches according to the color of the wires in the color hierarchy.

25 In general, a packet moves along a wire of a given color during a time frame if the packet seeks to move to a destination switch to which the wire is connected and if no other packet currently resides at the destination switch. However, if during a sending
30 switch step, a packet seeks to move along a wire from a source switch to a destination switch and another packet from a later initiated wave currently resides

- 9 -

at the destination switch, the position of the packets is swapped so that the packet previously at the source is at the destination and vice versa.

If, instead, a circuit switching strategy is
05 desired, the present invention provides means for extending paths in a circuit switching network. According to this method, a proposal is sent from a current node position in the network. For each message path that seeks extension, the proposal is
10 sent to each neighbor node in a desired direction of extension. Subsequent to the sending of a proposal, an acceptance is returned to the current node position from a neighbor node if the neighbor receives exactly one proposal. Upon acceptance, each
15 message path is advanced to include an accepting neighbor node if it has one.

To further enhance this method of extending message paths, the additional step of extending place-holders may be utilized. The place-holders are
20 sent on behalf of any message paths that are not moved forward during a given advancing step. The place-holders serve to reserve a place at a switch to which the message path seeks to extend. Thus, as the name implies the place-holders hold a place for a
25 stalled message path. Additionally, it is preferred that cancellation signals are sent from message path to place-holders when the place-holders are no longer needed. The cancellation signals result in the removal of the place-holders when the signals are
30 received. This prevents undue congestion within the switching network. It should be noted that the place-holders are advanced as if they are a message

-10-

path. Further, when a place-holder receives
cancellation signals from all message paths for which
it is holding a place, the place-holder sends a
cancellation signal to additional place-holders that
05 are reserves a spot for the place-holder.

Brief Description of the Drawings

Figure 1 depicts a 2x2 switch.

Figure 2 shows a splitter logical cluster.

Figure 3 illustrates a merger logical cluster.

10 Figure 4 shows a condenser logical cluster.

Figure 5 is an illustration of a butterfly
switching network.

Figure 6 depicts the blocks of a butterfly
network and a sample routing path.

15 Figure 7 depicts an example of a fat tree.

Figure 8 shows two butterfly networks and the
twin butterfly they form when merged.

Figure 9 shows an example twin multibutterfly
switching network.

20 Figure 10 shows an example 2-multi-Benes
switching network.

Figure 11 is an illustration of a splitter
having more than two levels.

25 Figure 12 is a flowchart of the major steps of
the packet switching method of the present invention.

Figure 13 is a flowchart of the edge coloring
scheme employed in a packet switching method of the
present invention.

30 Figure 14 is a flowchart of a method for
striking faulty input splitters in a switching
network to boost fault tolerance of the switching
network.

-11-

Figure 15 is a flowchart illustrating a method for removing faulty switches in a switching network to increase fault tolerance.

Figure 16 is a flowchart of the major steps of a
05 circuit switching algorithm.

Figure 17 depicts a non-blocking network.

Detailed Description of the Preferred Embodiment

Logical Clusters

10 The basic building block of many switching networks is the logical cluster. A logical cluster is a group of switches and wires that perform a high level task. Examples of logical clusters are
15 splitters, mergers and condensers. The splitter 10 shown in Figure 2 splits a set of inputs into two sets of outputs. Specifically, the splitter 10 has a set of inputs 4 that lead into a set of 2x2 switches 11 that constitute an input block 5. Wires 9 connect the input block 5 of switches to the output blocks 7a and 7b of switches 11. The outputs of these switches
20 11 are the two sets of outputs of the splitter 10. Each switch 11 in the input block 5 is connected to at least one output switch 11 in each of the output blocks 7a and 7b. The splitter 10 serves to route
25 the input 4 to the appropriate output block 7a and 7b. It does not matter which splitter output 6 a message is routed as long as each message is routed to a splitter output 6 in the correct output block 7a, 7b.

-12-

Another common variety of logical cluster is the merger 21 (Figure 3). The merger 21 is comprised of multiple input blocks such as 12a and 12b shown in Figure 3. In the merger 21 shown, the input blocks 12a and 12b lead via wires 9 into a single output block 13. Thus, the different sets of inputs feeding into the respective input blocks 12a and 12b are merged into a single set of outputs from output block 13. A third logical cluster that is used in switching networks is the condenser 23 (Figure 4). It condenses a set of inputs into a lesser number of outputs.

Butterfly Networks

A butterfly network is a common example of a switching network. It is referred to as a butterfly network because the connections between nodes form a pattern resembling a butterfly. A butterfly network has the same number of inputs as it has outputs. The inputs are connected to the outputs via a set of switches organized into successive levels of switches. An N-input, N-output butterfly network has $\log_2 N + 1$ (hereinafter \log_2 will be referred to as lg) levels of switches, each level having N 2x2 switches. For a message to travel from input to output, it must traverse at least one switch in each successive level.

An example butterfly network 8 is shown in Figure 2. Each switch 3 in the butterfly 8 has a distinct reference label $\langle L, r \rangle$, where L is its level, and r is its row. In an N-input butterfly, the level L is an integer between 0 and $\lg N$, and the row r is a $\lg N$ -bit binary number. The inputs and outputs reside

-13-

on levels 0 and $\lg N$, respectively. For $L < \lg N$, a switch labeled $\langle L, r \rangle$ is connected to switches $\langle L+1, r \rangle$ and $\langle L+1, r^{(L)} \rangle$, where $r^{(L)}$ denotes r with the L th bit complemented.

05 Butterfly networks are composed of sequences of splitters. The switches on each level of a butterfly network are partitioned into blocks according to which outputs they can reach. Another example butterfly network 19 is shown in Figure 6. The first
10 level of the butterfly network 19 can be viewed as a single block 7, since all of the inputs of the block 37 can reach all of the outputs of the block 37. The second level has two blocks: one block 14a consisting of those switches that can reach outputs whose labels
15 start with 0, and the other block 14b (shown in phantom form) consisting of those outputs whose labels start with 1. Each block 37, 14a, 14b, is the input group of a subsequent splitter having two output groups. Other sample blocks for the higher
20 levels include blocks 15a-15d and 17a-17h. Any pair consisting of an input and output of the butterfly is connected by a single logical (up-down) path through the butterfly. An example of such a logical path through a butterfly network is shown by the solid
25 lines 29 in Figure 4, indicating the decreasing number of switches that a message may choose from in the successive output blocks of the butterfly.

Fat-Tree

30 A fat-tree is another common example of a switching network that is made of splitters and mergers. A fat-tree network 16 is shown in Figure 7.

-14-

Its underlying structure is a complete 4-ary tree (i.e., every vertex has four wires leading to the next level of the tree). Each edge in the 4-ary tree corresponds to a pair of oppositely directed groups of wires called channels. The channel directed from the leaves 18 to the root 20 is called an up channel; the other channel is called a down channel. A group of up channels connecting four children to their parent forms a merger, while a group of down channels connecting a parent to its four children forms a splitter. A message routes up through the mergers until it can move down through the splitters to its destination. The capacity of a channel is the number of wires in the channel. The tree is referred to as "fat" because the capacities of the channels grow by a factor of 2 at every level. A fat-tree of height m has $M^2 - 2^{2m}$ leaves and $M - 2^m$ vertices at the root.

Degree

A logical cluster of a switching networks is said to be low-degree if the degree of the switches in the logical cluster is a small fixed constant (e.g., 4, 8, or 16) that is independent of the number of inputs or outputs in the logical cluster. A switching network is said to be low-degree if the degree of the switches in the network is a small fixed constant independent of the number of inputs or outputs in the network. For practical implementation of these networks, the low-degree property is a requirement due to fixed component pinout.

-15-

Expansion

The present invention is concerned with low degree logical clusters exhibiting expansion and/or dispersion. An N -input logical cluster is said to be

05 (α, β) -expansive if for every $k < \alpha N$, every set of k input ports collectively is connected to at least βk output ports in each group of outputs, where $\beta > 1$ and β and α are fixed threshold parameters independent of N . Typical values of α and β are $\alpha = 1/3$ and $\beta = 4/3$.

10 More simply, a logical cluster is said to be expansive if it is (α, β) -expansive for some threshold values of $\beta > 1$ and $\alpha < 1/\beta$. Intuitively, the expansion property for a logical cluster implies that for all

15 subsets of size k of the inputs to the cluster, there are more than k output switches in each output group which are connected to the subsets of size k . The possible paths that a message may assume, thus, expand between input and output. The extent to which there are more than k switches in the output group is

20 determined by the parameter β . Hence, if β is 2, there are twice as many output switches in each output group that are connected to the k inputs. k has a value less than the total number of inputs into the cluster N . The extent to which k is less than N

25 is determined by α .

A switching network comprised of expansive logical clusters is said to be expansive. Two examples of expansive switching networks are the multibutterfly network and the multi-Benes network,

30 both of which will be discussed below.

-16-

Multibutterfly

A multibutterfly network is formed by merging butterfly networks in a somewhat unusual manner. In particular, given 2 N-input butterflies G_1 and G_2 and
 05 a collection of permutations $(\Pi = \pi_0, \pi_1, \dots, \pi_{\lg N})$ where $\pi_L: [0, (N/2^L - 1)] \rightarrow [0, (N/2^L - 1)]$, a 2-butterfly is formed by merging the switch in row $(jN/2^L) + i$ of level L of G_1 with the switch in row $(jN/2^L) + \pi_L(i)$ of level L of G_2 for all $0 \leq i \leq (N/2^L) - 1$, all $0 \leq j <$
 10 $(2^L - 1)$, and all $0 \leq L \leq \lg N$. The result is a 2-butterfly comprising an N-input $\lg N + 1$ -level graph.

In other words, a twin multibutterfly 55 is formed from merging two butterflies such as 51 and 53 in Figure 8, each having N inputs. How the
 15 butterflies 51 and 53 are merged is determined by the permutation Π . In particular, given a first switch in the first butterfly located on level L at a given row (i.e., row = $jN/2^L + i$), the first switch is merged with a second switch in the second butterfly also on
 20 level L, but at a row determined by the permutation (i.e., $jN/2^L + \pi_L(i)$). The permutation maps an integer in the range of 0 to $N/2^L - 1$ to a permuted integer also in the range of 0 to $N/2^L - 1$.

For an example of such a butterfly, see the
 25 multibutterfly 55 shown in an enlarged view in Figure 9. Of the 4 output wires at a switch in the multibutterfly 55, two are up outputs and two are down outputs (with one up wire and one down wire being contributed from each butterfly). Thus, switch
 30 31 has two up wires 33 and two down wires 35. Multibutterflies (i.e. d-butterflies) are composed from d butterflies in a similar fashion using d-1 sets of permutations $\Pi^{(1)}, \dots, \Pi^{(d-1)}$, to produce a $\lg N$ level networks with $2d \times 2d$ switches.

-17-

The notion of up and down edges (or wires) discussed above relative to a 2-butterfly can be formalized in terms of splitters. More precisely, the edges from level L to level $L+1$ in rows $(jN/2^L)$ to $((j+1)N/2^L)-1$ in a multibutterfly form a splitter for all $0 \leq L < \lg N$ and $0 \leq j \leq 2^L-1$. Each of the 2^L splitters starting at level L has $N/2^L$ inputs and outputs. The outputs on level $L+1$ are naturally divided into $N/2^{L+1}$ up outputs and $N/2^{L+1}$ down outputs. All splitters on the same level L are isomorphic (i.e. they have the same number of inputs, the same number of outputs and the same wire connection patterns), and each input is connected to d up outputs and d down outputs according to the butterfly and the permutations $\pi_L^{(1)}, \dots, \pi_L^{(d-1)}$. Hence, any input and output of the multibutterfly are connected by a single logical (up-down) path through the multibutterfly, but each step of the logical path can be taken on any one of d edges.

An important characteristic of a multibutterfly is the set of permutations $\Pi^{(1)}, \dots, \Pi^{(d-1)}$ that prescribe the way in which the component butterflies are merged. For example, if all of the permutations are the identity map, then the result is the dilated butterfly (i.e., a butterfly with d copies of each edge).

Of particular interest to the present invention are multibutterflies that have expansion properties. A multibutterfly has expansion property (α, β) if each of its component splitters has expansion property (α, β) . In turn, an M -input splitter has expansion property (α, β) if every set of $k \leq \alpha M$ inputs is connected to at least βk up outputs and βk down outputs for $\beta > 1$.

-18-

If the permutations $\Pi^{(1)}, \dots, \Pi^{(d-1)}$ are chosen randomly, then there is a good probability that the resulting d-butterfly has expansion property (α, β) for any d, α , and β for which $2\alpha\beta < 1$ and

$$d < \beta + 1 + (\beta + 1 + \ln 2\beta) / \ln(1/2\alpha\beta) \quad (1)$$

It is not difficult to see that a multibutterfly network offers many advantages over a butterfly network. For example, whereas a butterfly contains just one path from each input port to each output port, a multibutterfly contains many paths from each input to each output port. Indeed, there is still just one logical (up-down) path from any input to any output, but this logical path can be realized as any one of several physical paths.

Another advantage of a multibutterfly (or any switching network comprised of expansive logical clusters) is that it is very hard for a large number of switches to become blocked by congestion or by faults. The reason is that for k inputs of an expansive logical cluster to become blocked, at least βk of the outputs would have to be blocked or faulty themselves. In a typical network like the butterfly, the reverse is true. One can block k inputs by congesting only k/2 outputs. When this property is cascaded over several levels of the network, the difference in performance between the butterfly versus the multibutterfly can be dramatic. For example, one fault can block 1000 switches 10 levels back in the butterfly, but 1000 faults would be necessary to block just 1 switch 10 levels back in the multibutterfly (if $\beta=2$).

-19-

Multi-Benes

Like a multibutterfly, a multi-Benes network is formed from merging networks together, specifically Benes networks. A 2-multi-Benes network 26 is shown in Figure 9. An N-input multi-Benes network has $2\lg N + 1$ levels labeled 0 through $2\lg N$. Levels $\lg N$ through $2\lg N$ form a multibutterfly, while levels 0 through $\lg N$ form the mirror image of a multibutterfly. Thus, informally the 2 multi-Benes network can be viewed as a network made of two butterfly networks placed back to back.

As in the multibutterfly, the edges in levels $\lg N$ through $2\lg N$ of the multi-Benes are partitioned into splitters. Between levels 0 and $\lg N$, however, the edges are partitioned into mergers. More precisely, the edges from level L to level L+1 in rows $j2^{L+1}$ to $(j+1)2^{L+1} - 1$ form a merger for all $0 \leq L < \lg N$ and $0 \leq j \leq N/2^{L+1} - 1$. Each of the $N/2^{L+1}$ mergers starting at level L has 2^{L+1} inputs and outputs. The inputs on level L are naturally divided into 2^L up inputs and 2^L down inputs. All mergers on the same level L are isomorphic, and each input is connected to 2^L outputs. There is a single (trivial) logical path from any input of a multi-Benes network through the mergers on the first $\lg N$ levels to the single splitter on level $\lg N$. From level $\lg N$, there is a single logical path through the splitters to any output. In both cases, the logical path can be realized by many physical paths.

An M-output merger has expansion property (α, β) if every set of $k \leq \alpha M$ inputs (up or down) is connected to at least $2\beta k$ outputs where $\beta > 1$. With nonzero probability, a random set of permutations

- 20 -

yields a merger with expansion property (α, β) for any d, α , and β for which $\alpha\beta < 1$ and

$$2d < 2\beta + 1 + (2\beta + 1 + \ln 2\beta) / \ln(1/2\alpha\beta) \quad (2)$$

05 A multi-Benes network has expansion property (α, β) if each of its component mergers and splitters has expansion property (α, β) . The multibutterflies and multi-Benes networks considered herein are assumed to have expansion property (α, β) unless otherwise stated.

10 Dispersive Logical Clusters

The present invention is also concerned with dispersive logical clusters. An N -input logical cluster is said to be dispersive if for prespecified threshold parameters α, δ , for every $k \leq \alpha N$ and for
15 every set of k inputs, there are at least δk outputs ports in each group of the logical cluster that are connected to precisely one of the k input ports. This definition of a dispersive logical cluster may be restated as: a logical cluster wherein every
20 subset x of $k \leq \alpha N$ inputs in the cluster, there are δk switches in x which have a neighbor in each output group that is not connected to any other switch in x . In other words, δk nodes in x have a unique neighbor for each output group. This property is called the
25 unique-neighbors property in S. Arora, T. Leighton and B. Maggs, On-line algorithms for path selection in a non-blocking network, Proceedings of the 22nd Annual ACM Symposium on the Theory of Computing, 1990. A switching network comprised of dispersive
30 logical clusters is called a dispersive switching network.

-21-

A switching network with dispersive logical clusters offers substantial advantages over ordinary switching networks. Ordinary networks are typically unable to concurrently move all messages forward. In most instances, some messages are moved forward while others are terminated or queued for long periods of time. In a network with dispersive logical clusters, on the other hand, a message at an input can advance without fear of blocking other messages if the input is connected to an output which is coupled to no other inputs with a message. Given that in a dispersive logical cluster at least δk outputs in each output group of a cluster are connected to precisely one input switch, at least a δ fraction of all messages can advance at every step in a dispersive logical cluster without ever blocking other messages.

Any splitter with the (α, β) expansion property has the (α, δ) dispersion property where $\delta = 2\beta/d - 1$, provided that $\beta > d/2$. See Arora et al., supra. By Equation 1, it is evident that randomly generated splitters have the (α, δ) dispersion property where δ approaches 1, as d gets large and as α gets small. Explicit constructions of such splitters are not known, however. Only multibutterflies with the (α, δ) dispersion property for $\delta > 0$ will be discussed below. It should be noted that the (α, β) expansion property (where $\beta > d/2$) is a sufficient condition for the dispersion property, but by no means necessary. In fact, the existence of random splitters which have a fairly strong (α, δ) dispersion property for small degree is proven in Arora et al.

-22-

Amongst the many methods for constructing expansive and/or dispersive logical clusters with low-degree is a method consisting of connecting the input ports to the output ports randomly so that every switch has the same degree. With high probability, the resulting logical cluster is expansive and dispersive if the number of connections from each input port to each output port is 3 or higher. Even if there are only 2 connections from each input port to each output port, the concatenation of two logical clusters still has the expansive and dispersive properties with high probability.

Logical clusters such as splitters and mergers have been described thus far as only having a depth of 1 (i.e., input ports are switches that are directly connected to output ports). Nevertheless, the logical clusters can have more than one level. For example, if two depth 1 splitters are cascaded, a depth $2N$ splitter is generated (See Figure 11). Logical clusters with depth 1, 2, or 3 are of primary interest in the present invention. For logical clusters with depth 2 or more, an input port is connected to an output port if it can be connected to the output port by an appropriate setting of the switches in the logical cluster. The definitions of expansion and dispersion given above apply equally to such logical clusters. Using these definitions, it is possible to construct depth 2 logical clusters comprised of 2×2 switches that have the expansion and the dispersion property (see Arora, supra; and T. Leighton and B. Maggs, Expanders might be practical:

- 23 -

fast algorithms for routing around faults in multi-butterflies, Proceedings of the 30th Annual Symposium on Foundations of Computer Science, October 1989, pps. 384-389).

05

Routing Methods

In addition to the new types of switching networks, the present invention is concerned with methods for routing messages on such networks. These routing methods allow many messages to be routed to
10 their correct destination quickly using only destination addresses and on-line control. On-line control refers to all decisions regarding switching being made locally by each switch without global information about the location and destinations of
15 other packets. The routing methods also allow the messages to be routed around faulty switches and/or busy switches. A switch is said to be faulty if it is not functioning correctly, and it is said to be busy if it cannot be used to route any additional
20 messages (i.e., all of its capacity is currently being used).

These features of the routing methods can be integrated into the notion of switch availability. In particular, the status of an input switch in a
25 logical cluster is "available" if all of the following conditions are met:

1. the switch is not faulty;
2. the switch is not busy; and

-24-

3. the switch is connected to at least δ available output ports in each output group, where δ is a prespecified threshold value that is at least one.
- 05 Otherwise, the switch is said to be "unavailable".
By only sending messages to "available" switches, the routing methods are able to avoid switches on multibutterflies that are faulty or busy. Such an approach also avoids the routing of a message into a
10 position wherein the message might only be further routed to a faulty or busy switch.

Packet Switching

- It is possible to run a variety of package switching methods, e.g., a greedy algorithm.
- 15 Following is a preferred greedy algorithm. In describing the preferred packet switching method, it is assumed unless stated otherwise that the multibutterfly networks being used have expansion property (α, β) for $2\alpha\beta < 1$ and $\beta > 1$. A flowchart
20 outlining the major steps of the packet switching method is shown in Figure 12. Initially, packets to be routed across the switching network are partitioned into waves (Step 50) so that at most one packet in each wave is destined for any set of Z
25 contiguous outputs. One way to achieve such a partitioning into waves is to group packets into the same wave if they are in the same permutation and their destinations are congruent modulo Z . For P permutations to be routed, this approach of
30 partitioning results in at most PL waves. In

-25-

- general, Z should be set to equal $1/(2\alpha)$, since then it is certain that at most $M/(2Z) = \alpha M$ packets in any wave pass through the up (or down) edges of any M -input splitter of the multibutterfly (for any M).
- 05 This allows the (α, β) expansion property to apply to the set of inputs of any splitter occupied by the packets of a single wave at any time. (E.g., if k inputs of a splitter contain packets of a single wave that want to traverse up edges, then these inputs are
- 10 connected to at least βk up outputs.) This is because packets going through the $M/2$ up (or $M/2$ down) splitter outputs can only be destined for the descendant set of $M/2$ contiguous multibutterfly outputs.
- 15 The routing of the packets proceeds in stages (see Steps 52 and 54), wherein each stage consists of an even phase (Step 56) and an odd phase (Step 58), and each phase consists of $2d$ steps. In even phases, packets are sent from even levels to the next odd
- 20 levels (Step 56), and in odd phases, packets are sent from the odd levels to the next even levels (Step 58).
- The edges connecting levels are colored in $2d$ colors so that each node is incident to one edge of each color. In each phase, the edges are processed
- 25 by color in sequence for all colors such that one step is dedicated per color. A flowchart of the activity performed in a step is provided in Figure 3. For each step (see Step 62), a packet is moved
- 30 forward along an edge with the color being moved during the step (Step 68) provided that there is a packet in the switch at the tail of the edge that

-26-

wants to go in that direction (up or down) (Step 64) and further provided that there is no packet in the switch at the head of the edge (Step 66).
Alternatively, if there is a packet in the switch at
05 the head of the edge (Step 66) and if the packet is in a later wave than the packet at the tail of the edge (Step 70), the two packets are swapped (Step 72) so that the packet in the earlier wave moves forward. Otherwise, the packet is not moved (Step 74). Note
10 that every switch processes and/or contains at most one packet at any step.

If there is only one permutation to route, then each input of the multibutterfly starts with one packet. If there are P permutations to be routed,
15 however, it is useful to augment the front-end of the multibutterfly with $P - 1$ levels of d (random) matchings so that the queue size of 1 at the input level can be preserved. The augmentation requires no more hardware than that necessary to augment the
20 front end of each component butterfly with a $P - 1$ cell linear array. Moreover, the augmentation ensures that the preprocessing levels have an (α, β) -expansion property at least as strong as the first level. For notational purposes, these
25 additional levels will be referred to hereinafter as levels $-1, -2, \dots, -(P-1)$.

The waves, edge coloring, and odd-even phases can be dispersed for most applications. Specifically, each packet is forwarded unless all
30 queues ahead of it exceed some prespecified threshold of fullness. This approach is denoted as the greedy algorithm. For more details on the greedy algorithm, see Arora, supra; and Leighton, supra.

-27-

Fault Tolerance

The present invention also embodies a method for fault tolerance in networks such as the multibutterfly. The central idea of the method is to identify and remove those parts of the network that contain too many faulty switches to be used. The goal of this reconfiguration process is to salvage as much of the working hardware as possible while leaving largely intact the expansion property of the network. Once an appropriate set of inputs and outputs have been removed, the greedy algorithm described in the previous section can be applied to route packets between the remaining inputs and outputs.

The first step in the fault tolerance method is to specify the outputs to remove. A flowchart of the output removal scheme is given in Figure 14. In particular, each splitter in the multibutterfly is examined (Steps 80 and 82). If more than an ϵ_0 fraction of the input switches are faulty (Step 84), where $\epsilon_0 = 2\alpha(\beta' - 1)$ and $\beta' = \beta - (d/2)$, then the splitter is "erased" from the network (Step 86). In addition all of its descendant switches and outputs are likewise "erased" (Step 86).

The fault tolerance method next involves determining which inputs to remove (see flowchart in Fig. 15). The first step (87) in the process is to declare any faulty switch as unavailable. Working from the $(\lg N)$ th output level backwards (see Steps 89, 90, 100, and 102), each switch is examined (Box 90) to determine if at least half of its upper outputs lead to faulty unavailable that have not been erased (Step 92), or if at least half of its lower outputs lead to unavailable switches that have not been

-28-

erased (Step 94). If so, the switch is declared
(Step 96) to be unavailable. (But it is not erased.
Where outputs lead to erased switches, they need not
be declared unavailable in subsequent checking of
05 preceding levels because the outputs from the erased
switches are invalid). This process is repeated for
all switches on a level (see Step 98) and for each
level (Steps 98 and 100) until all levels have been
examined (Step 102).

10 All the remaining unavailable switches are
erased (Step 104). What is left is a network in
which every input in every splitter is linked to $d/2$
functioning upper outputs (if the descendant
multibutterfly outputs exist) and $d/2$ functioning
15 lower outputs (if the corresponding multibutterfly
outputs exist). Hence, every splitter has an (α, β')
expansion property. Thus, it can be proven that the
greedy algorithm still routes any permutation on the
remaining inputs and outputs quickly. See Arora,
20 supra; and Leighton, supra for more details.

Circuit Switching

The above described routing method works well
for packet switching. It is not as appropriate for
direct application to circuit switching.

25 Nevertheless, a similar approach can be used for
circuit switching. The circuit-switching method
adopts an approach similar to the packet switching
method but additionally makes use of the dispersion
property of the logical clusters.

30 In order for the circuit-switching algorithm to
succeed, the multibutterfly network must be lightly

-29-

loaded by some fixed constant factor Z . Thus, in an N -row multibutterfly network, connections are made only between the N/Z inputs and outputs in rows that are multiples of Z . Since the other inputs and
05 outputs are not used, the first and last $\lg Z$ levels of the network can be removed, and the N/Z inputs and outputs can each be connected directly to their Z descendants and ancestors on levels $\lg Z$ and $\lg N - \lg Z$, respectively.

10 It is relatively easy to extend paths from one level to the next in a multibutterfly with the (α, δ) dispersion property. The reason is that those paths at switches with unique neighbors can be trivially extended without worrying about blocking any other
15 path trying to reach the next level. By proceeding recursively, it is easy to see that all the paths can be extended from level L to level $L+1$ (for any L) in $\log(N/Z2^L)/\log(1/L-\delta)$ steps. When such a recursive approach is adopted, the method proceeds in steps
20 wherein each "step" consists of (see flowchart in Figure 16):

1. sending out a "proposal" for every path still waiting to be extended to the output (level $L+1$) neighbors in the desired
25 direction (up or down) (Step 106);
2. sending back acceptance of the proposal from every output node that receives precisely one proposal (Step 108). If more than one proposal is received, none are
30 accepted;
3. advancing every path receiving an acceptance to one of its accepting outputs on level $L+1$ (Step 110).

-30-

Splitters connecting level L to level $L+1$ have $M = N/2^L$ inputs. At most M/Z paths can pass through splitters connecting level L and level $L+1$ by definition of Z . Since $Z > 1/\alpha$, the set of switches
 05 containing paths needing to be extended has a size of at most αM . The (α, δ) dispersion property can be applied to ensure that at each step, the number of paths still remaining to be extended decreases by a $(1-\delta)$ factor. Hence, all of the paths are extended
 10 in $\log(N/Z2^L)/\log(1/(1-\delta))$ steps, as claimed.

By using the path extension algorithm just described on each level in sequence, all of the paths can be constructed in

$$15 \quad \sum_{l=0}^{\lg N - 1} \frac{\log \frac{N}{2^l}}{\log \frac{1}{1-\delta}} \leq \frac{\log^2 \frac{N}{2^l}}{2 \log \frac{1}{1-\delta}} = O(\log^2 N)$$

bit-steps. To construct the paths in $O(\log N)$ bit-steps the algorithm may be modified. Specifically, given a fraction $< \alpha$ of paths that need
 20 to be extended at an M -input splitter, the method does not wait $O(\log M)$ time for every path to be extended before it begins the extension at the next level. Instead, it waits only $O(1)$ steps, in which time the number of unextended paths falls to a
 25 fraction ρ of its original value, where $\rho < 1/d$. Now the path extension process can start at the next level. The danger here is that the ρ fraction of paths left behind may find themselves blocked by the time they reach the next level, and so it is
 30 necessary to ensure that this will not happen. Therefore, stalled paths send out place-holders to all of their neighbors at the next level, and henceforth the neighbors with place-holders

-31-

participate in path extension at the next level as if the place-holders were paths. Of course, the neighbors holding place-holders must in general extend in both the upper and the lower output
05 portions of the splitter, since they do not know yet which path will ultimately use them.

It is worth noting that a place-holder not only reserves a spot that may be used by a path at a future time, but also helps to chart out the path by
10 continuing to extend ahead. In order to prevent place-holders from multiplying too rapidly and clogging the system (since if the fraction of inputs of a splitter which are trying to extend rises above α , the path extension algorithm may cease to work),
15 it is necessary to ensure that as stalled paths get extended, they send cancellation signals to the place-holder nodes ahead of them to indicate that the place-holder nodes are no longer needed. When a place-holder node receives cancellation signals from
20 all the nodes for which the place-holder node was holding a place, the place-holder node is removed and ceases to extend anymore. In addition, the place-holder nodes send cancellations to any nodes ahead of them that may be holding a place for them.

25 The $O(\log N)$ -step algorithm for routing paths proceeds in phases consisting of the following two types of steps:

1. Steps of passing cancellation signals.
There are C such steps. The cancellation
30 signals travel at the rate of one level per step;

-32-

2. Steps of extending from one level to the next. There are T such steps. In this time, the number of stalled (i.e., unextended) paths at each splitter drops by at least a factor of ρ , where $\rho \leq (1-\delta)^T$.

Each path is restricted to extend forward by at most one level during each phase. The first wave of paths and place-holders to arrive at a level is referred to as the wavefront. The wavefront moves forward by one level during each phase. If a path or a place-holder in the wavefront is not extended in T steps, it sends place-holders to all of its neighbors at the end of the phase. It is assumed that $C \geq 2$ so that cancellation signals have a chance to catch up with the wavefront. It is also assumed that $d \geq 3$. (See Arora, supra for more details and a proof that the algorithm quickly establishes all paths).

Nonblocking Networks

The preceding method works well for static message routing problems (i.e., for problems in which the messages to be routed all start at the input ports of the circuit at the same time). In some applications, however, such as telephone networks, the messages to be routed arrive at the input ports at different times. In such cases, it is desirable to route the messages in a nonblocking manner. The goal of such routing switches is to interconnect the terminals and switches so that any unused input-output pair can be connected by a path of unused switches, regardless of the other paths that exist at the time. Such a network is said to be

-33-

nonblocking. Nonblocking in this strong sense is to be distinguished from the rearrangeable properties of Benes networks which allow further connections but require rerouting of existing connections. The
05 6-terminal graph 41 shown in Figure 17 is nonblocking since no matter which input-output pairs are connected by a path, there is a node-disjoint path linking any unused input-output pair. As such, if Bob is talking to Alice and Ted is talking to Carol,
10 then Pat can still call Vanna.

To satisfy connection requests in a non-blocking multi-Benes network, an algorithm is employed that establishes a path from an unused input to an unused output in $O(\log N)$ bit-steps, where N is the number of
15 rows in the network. In the description of the algorithm that follows, it is assumed that at most one input tries to access any output at a time, and that each input accesses at most one output at a time.

20 The central idea behind the nonblocking routing algorithm used for an expansive switching network is to treat the switches through which paths have already been established as if they were faulty and to apply the previously described status propagation
25 techniques to the network. In particular, a node is defined as busy if there is a path currently routing through it. At the start, all busy nodes are said to be unavailable. A node is defined recursively to be unavailable if all of its up outputs or if all of its
30 down outputs are unavailable. More precisely, switches are declared to be unavailable according to the following rule. Working backwards from level $2\lg N - \lg 2 - 1$ to level $\lg N$, a switch is declared

-34-

unavailable if either all d of its up edges or all d of its down edges lead to unavailable switches. From level $\lg N - 1$ to level $\lg Z$, a switch is declared unavailable if all $2d$ of its outputs lead to
05 unavailable switches. A switch that is not unavailable is said to be available.

After the status propagation process, every available switch in the first half of the network has an output that leads to an available switch, and
10 every available switch in the second half has both an up output and a down output that leads to available switches. Furthermore, since at most a 2α fraction of the switches in each merger on level $\lg Z$ are unavailable, each of the N/Z inputs has an edge to an
15 available switch on level Z . At the other end, each of the N/Z outputs can be reached by an available switch on level $2\lg N - \lg Z$. As a consequence, a path can be established through available switches from any unused input to any unused output in $O(\log N)$
20 bit-steps using a simple greedy algorithm. Since the declaration of unavailable switches takes just $O(\log N)$ bit-steps, and since the greedy routing algorithm is easily accomplished in $O(\log N)$ bit-steps, the entire process takes just $O(\log N)$
25 bit-steps.

It is not difficult to implement the circuit-switching algorithm for use with a multi-Benes network. Specifically, the definition of being blocked must be modified so that a node on level L is
30 blocked if more than $2\beta - d - 1$ of its up (or down) neighbors on level $L+1$ are unavailable. (As before, it is assumed that $\beta > d/2$.) Available nodes are then

-35-

guaranteed to have at least $2d - 2\beta + 1$ available neighbors. Hence, any set of $k \leq \alpha M$ available inputs in an M-input splitter has a $(\alpha, 1/d)$ dispersion property, which is sufficient for the routing algorithm to work. Of course, it must also check that the modified definition of unavailable does not result in any inputs to the multi-Benes network becoming unavailable. See Arora, supra and Leighton, supra for details.

10 The above process can be applied in a fault tolerant environment by initially declaring all faulty switches as unavailable.

Although preferred embodiments have been specifically described and illustrated herein, it will be appreciated that many modifications and variations of the present invention are possible, in light of the above teachings, within the purview of the following claims, without departing from the spirit and scope of the invention.

- 36 -

CLAIMS

1. In a switching network for routing messages, a logical cluster, comprising:
 - 05 a) first and second sets of switches having N inputs for receiving messages and also having N outputs for outputting messages, wherein the second set of switches is divided into at least two unique groups of switches;
 - 10 b) connectors for connecting the first set of switches to the second set of switches such that the logical cluster exhibits a dispersion property wherein for every set of k inputs to the switches in the first set of switches, there are at least δk outputs in each
15 group of the second set of switches that are connected to precisely one of the k inputs to the first set of switches, where $k \leq \alpha N$ and k and N are positive integers, where δ and α are positive constants less than one.
- 20 2. A switching network as recited in Claim 1 wherein the switching network is a multibutterfly switching network.
- 25 3. A switching network as recited in Claim 1 wherein the switching network is a multi-Benes switching network.

-37-

4. A switching network as claimed in Claim 1 organized into levels of switches comprised of logical clusters of switches that are interconnected by connections such that, for each logical cluster of switches, connections connect each switch of an input group of switches to an output switch in each respective output group of switches, the network further comprising means for declaring unusable switches as unavailable and, proceeding from an output level backward, declaring each switch unavailable if the switch does not have a sufficient quantity of connections to non-unavailable switches in each output group for each logical cluster, and means for avoiding switches declared unavailable in routing messages across the switching network.
5. A network as claimed in Claim 4 wherein a busy switch or a faulty switch is unusable.
6. A network as claimed in Claim 4 further comprising means for deactivating all switches in a logical cluster as well as all descendant switches of the logical cluster where a number of faulty input switches of that logical cluster exceeds a predetermined threshold.

-38-

7. A network as claimed in Claim 1 further comprising means for extending message paths between nodes of the network comprising:
- 05 means for sending a proposal from a current node position for each message path that is to be extended to each neighbor node in a desired direction of extension;
- 10 means for returning an acceptance to the proposal to the current node position from a neighbor node if the neighbor receives exactly one proposal; and
- means for advancing each message path to include an accepting neighbor node for each message path receiving exactly one acceptance.
- 15 8. A network as claimed in Claim 7 further comprising means for sending placeholders on behalf of any message paths that are not moved forward such that the placeholders reserve a place at a switch to which the message path is
- 20 to extend.
9. A method of extending message paths in a circuit switching network having nodes, comprising the steps of:
- 25 a) sending a proposal from a current node position for each message path that is to be extended to each neighbor node in a desired direction of extension;

-39-

- b) returning an acceptance to the proposal to the current node position from a neighbor node if the neighbor receives exactly one proposal; and
 - 05 c) advancing each message path to include an accepting neighbor node for each message path receiving exactly one acceptance.
10. A method as recited in Claim 9 wherein the switching network is a multibutterfly network.
- 10 11. A method as recited in Claim 9 wherein the switching network is a multi-Benes network.
12. A method as recited in Claim 9 further comprising the step of sending place-holders on behalf of any message paths that are not moved
15 forward during the advancing step wherein said place-holders reserve a place at a switch to which the message path is to extend.
13. A method as recited in Claim 12 wherein during the advancing step, place-holders are advanced
20 as if they are message paths.
14. A method as recited in Claim 9 wherein cancellation signals are sent from message paths to place-holders when the place-holders are no longer needed, said cancellation signals
25 resulting in the removal of said place-holders when received.

-40-

15. A method as recited in Claim 14 wherein when a place-holder receives cancellation signals from all message paths for which it is holding a place, the place-holder sends a cancellation signal to additional place-holders that are holding a place for the place-holder.
16. In a multi-Benes switching network for routing messages, a low degree logical cluster, comprising:
- a) first and second sets of switches having inputs for receiving messages and outputs for outputting messages, wherein the second set of switches is divided into one or more disjoint groups of switches, each switch in the first and second sets of switches making local routing decisions; and
 - b) connectors for connecting the first set of switches to the second set of switches such that each output of the first set of switches is connected to an input of a switch in each of the groups of switches in the second set of switches and further such that the logical cluster exhibits an expansion property wherein there exists for every set of k switches in the first set of switches at least βk switches in each group of the second set of switches are connected to the k outputs of the first set of switches, where $\beta > 1$, and $k < \alpha N$ for N equal to the number of inputs into the first set of switches and for $\alpha < (1/\beta)$.

-41-

17. A multi-Benes switching network as claimed in Claim 16, comprising d Benes switching networks merged together and numbered 1 through d , where d is a positive integer, each network having
- 05 rows and levels of switches and N inputs, said networks being merged such that given a set of permutations $(\Pi^1, \dots, \Pi^{(d-1)})$ where $\Pi^k < \pi_0^k, \pi_1^k, \dots, \pi_{2^{lgN}^k}^k >$ and where $\pi_L^k: [0, N/2^{lgN-L-1}] \rightarrow [0, N/2^{lgN-L-1}]$ for $0 \leq L \leq$
- 10 lgN , a switch in row $(jN/2^{lgN-L}) + i$ of level L of Benes switching network number k is merged with a switch in row $(jN/2^{lgN-L}) + \pi_L^k(i)$ of level L of Benes switching network number $k+1$ for all $1 \leq k \leq (d-1)$, all $0 \leq i \leq (N/2^{lgN-L})-1$, all $0 \leq j \leq (2^{lgN-L})-1$, and all $0 \leq L \leq lgN$, and further
- 15 given $\pi_L^k: [0, N/2^{L-LgN-1}] \rightarrow [0, N/2^{L-LgN-1}]$ for $lgN \leq L \leq 2lgN$, a switch in row $(jN/2^{L-LgN}) + i$ of level L of Benes switching network number k is merged with a switch in row
- 20 $(jN/2^{L-LgN}) + \pi_L^k(i)$ for all $1 \leq k \leq (d-1)$, all $0 \leq i \leq (N/2^{L-LgN})-1$, all $0 \leq j \leq (2^{L-LgN})-1$ and all $lgN \leq L \leq 2lgN$.
18. A multi-Benes switching network as recited in Claim 17 wherein $d = 2$.
- 25 19. A multi-Benes switching network as recited in Claim 17 where permutations $(\Pi^1, \dots, \Pi^{(d-1)})$ are all the identity map.

-42-

20. A method of increasing fault tolerance of a switching network organized into levels of switches comprised of logical clusters of switches that are interconnected by connections such that, for each logical cluster of switches, connections connect each switch of an input group of switches to an output switch in each respective output group of switches, comprising the steps of:
- 05
- 10 a) proceeding from an output level backwards, declaring unusable switches as unavailable and examining output connections of switches and declaring a switch unavailable if the examining step reveals that the
- 15 switch does not have a sufficient quantity of connections to non-unavailable switches in each output group for each logical cluster; and
- 20 b) avoiding switches declared unavailable in routing messages across the switching network.
21. A method as recited in Claim 20 wherein the switching network is a multibutterfly switching network.
- 25 22. A method as recited in Claim 20 wherein the switching network is a multi-Benes network.
23. A method of claimed in Claim 20 wherein a busy switch is unusable.

-43-

24. A method as claimed in Claim 20 wherein a faulty switch is unusable.
25. A method as claimed in Claim 20 comprising the steps of:
- 05 (a) inspecting input switches of each logical cluster to determine how many of the input switches are faulty in said logical cluster; and
- 10 (b) where a number of faulty input switches of a logical cluster exceeds a predetermined threshold, deactivating all switches in the logical cluster as well as all descendant switches of the logical cluster so that they may not be used in routing messages
- 15 across the switching network.

1/10




FIG. 1 (PRIOR ART)




FIG. 2 (PRIOR ART)




FIG. 3 (PRIOR ART)




FIG. 4 (PRIOR ART)




FIG. 5 (PRIOR ART)




FIG. 6 (PRIOR ART)




FIG. 7 (PRIOR ART)




FIG. 8




FIG. 9


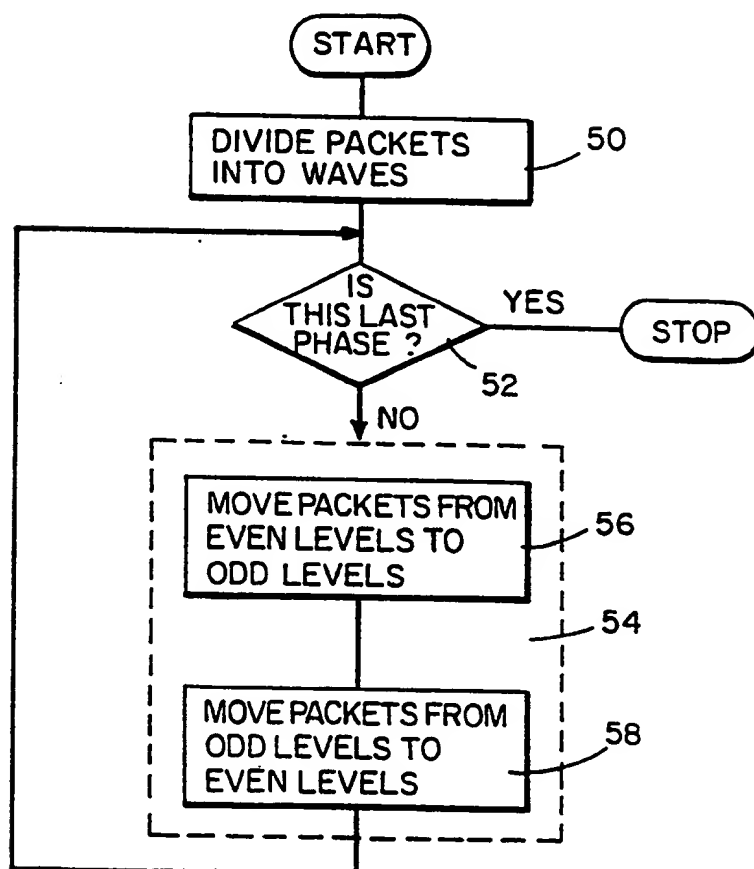



FIG. 10

6/10



7/10




FIG. 13

8/10

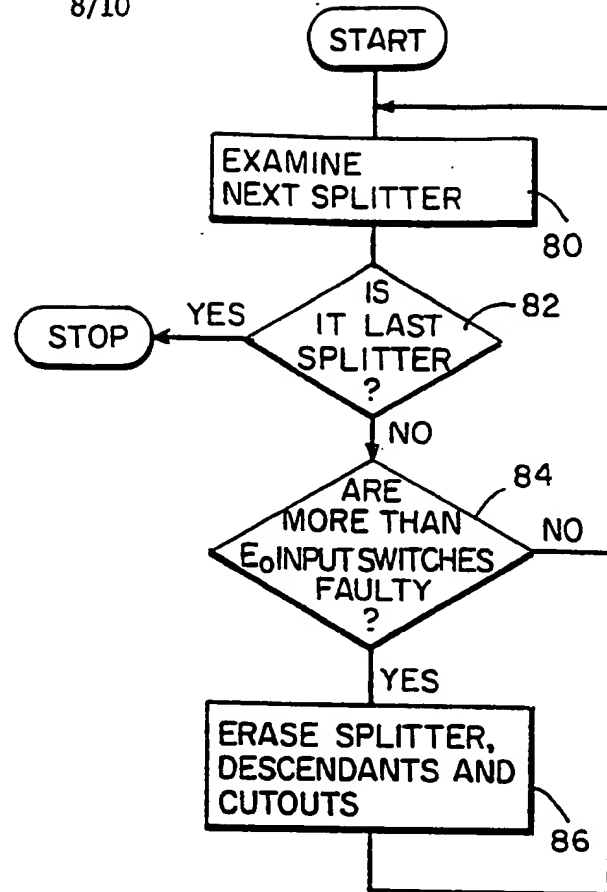


FIG. 14




FIG. 16

SUBSTITUTE SHEET

9/10

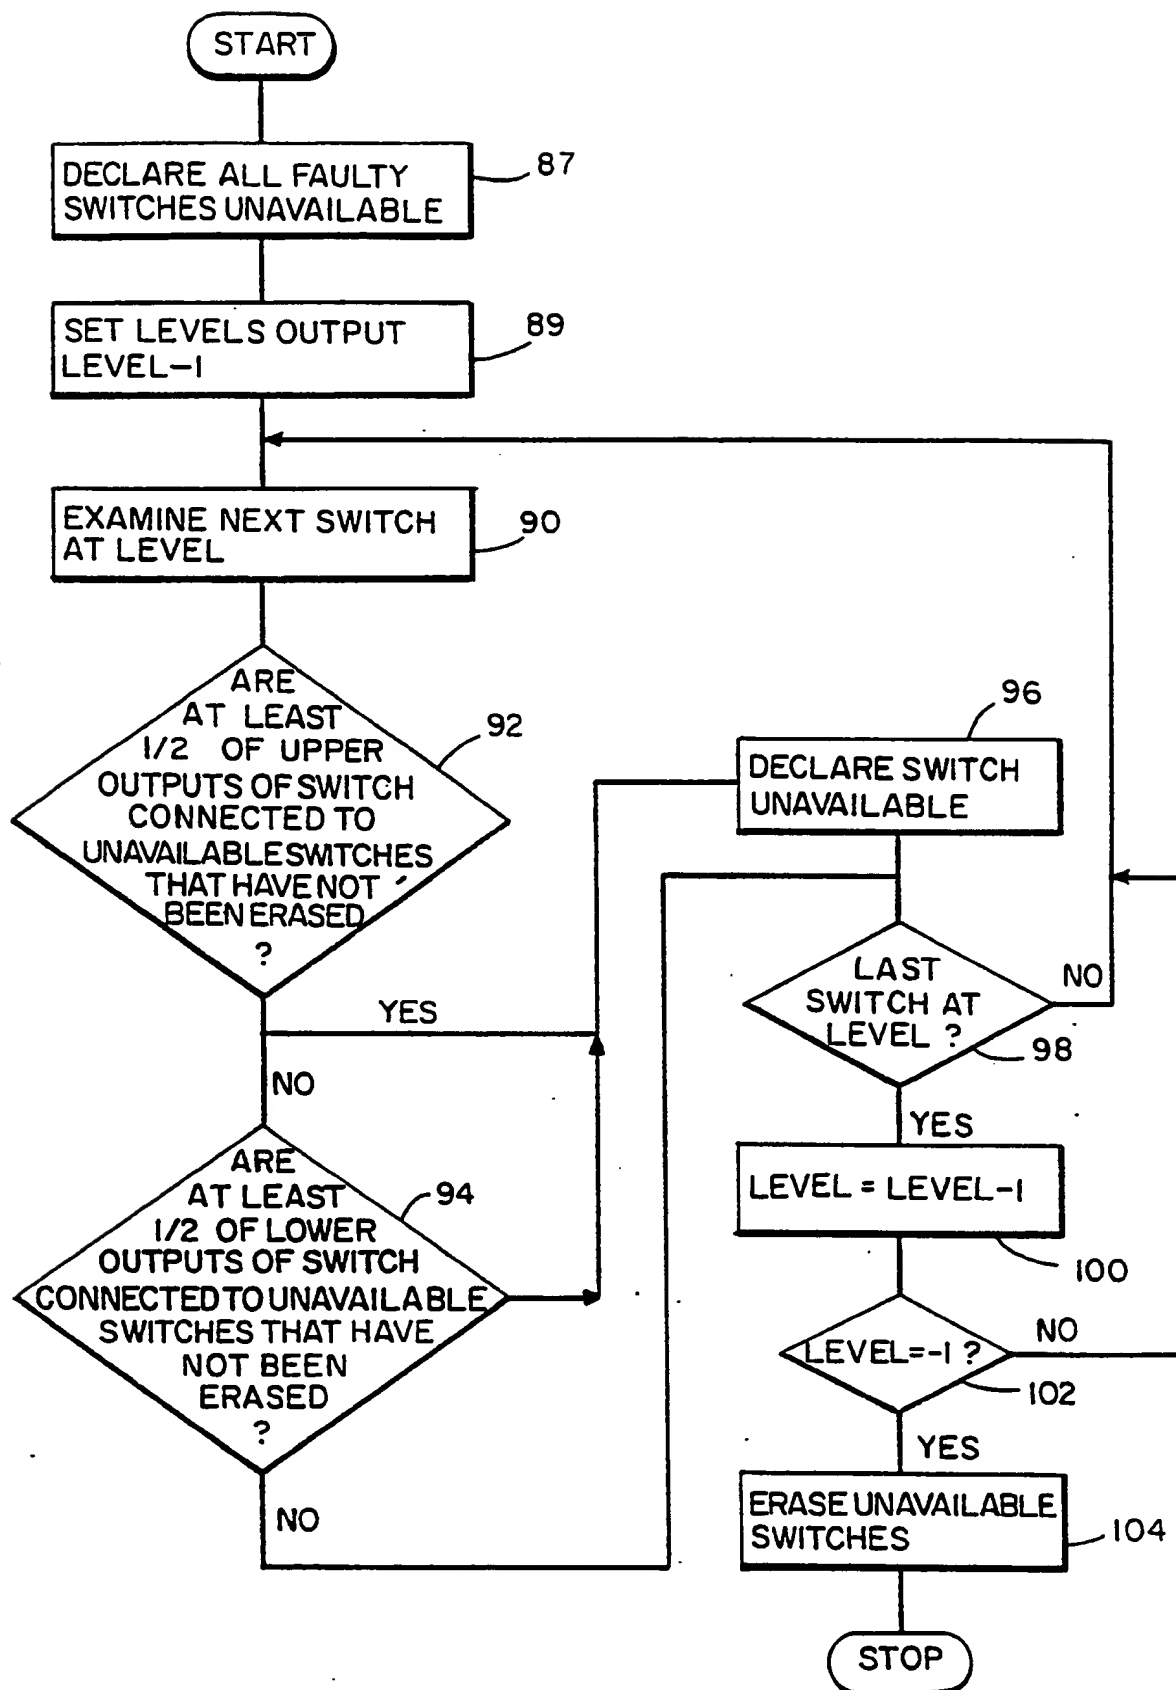


FIG. 15

SUBSTITUTE SHEET

10/10




FIG. 17